

## FREE: Faster and Better Data-Free Meta-Learning

Yongxian Wei<sup>1</sup> Zixuan Hu<sup>1</sup> Zhenyi Wang<sup>2</sup> Li Shen<sup>3,\*</sup> Chun Yuan<sup>1,\*</sup> Dacheng Tao<sup>4</sup>  
<sup>1</sup>Tsinghua Shenzhen International Graduate School; <sup>2</sup>University of Maryland, College Park  
<sup>3</sup>JD Explore Academy; <sup>4</sup>Nanyang Technological University

weiyx23@mails.tsinghua.edu.cn; huzixuan21@mails.tsinghua.edu.cn; zwang169@umd.edu  
 mathshenli@gmail.com; yuanc@sz.tsinghua.edu.cn; dacheng.tao@ntu.edu.sg

### Abstract

*Data-Free Meta-Learning (DFML) aims to extract knowledge from a collection of pre-trained models without requiring the original data, presenting practical benefits in contexts constrained by data privacy concerns. Current DFML methods primarily focus on the data recovery from these pre-trained models. However, they suffer from slow recovery speed and overlook gaps inherent in heterogeneous pre-trained models. In response to these challenges, we introduce the **Faster and Better Data-Free Meta-Learning (FREE)** framework, which contains: (i) a meta-generator for rapidly recovering training tasks from pre-trained models; and (ii) a meta-learner for generalizing to new unseen tasks. Specifically, within the module **Faster Inversion via Meta-Generator**, each pre-trained model is perceived as a distinct task. The meta-generator can rapidly adapt to a specific task in just five steps, significantly accelerating the data recovery. Furthermore, we propose **Better Generalization via Meta-Learner** and introduce an implicit gradient alignment algorithm to optimize the meta-learner. This is achieved as aligned gradient directions alleviate potential conflicts among tasks from heterogeneous pre-trained models. Empirical experiments on multiple benchmarks affirm the superiority of our approach, marking a notable speed-up (20×) and performance enhancement (1.42% ~ 4.78%) in comparison to the state-of-the-art.*

### 1. Introduction

Data-Free Meta-Learning (DFML) [16, 17, 19, 40] aims to derive knowledge from a collection of pre-trained models without necessitating the original data, enabling the adaptation of knowledge to new unseen tasks. Traditional meta-learning methods assume access to a collection of tasks with available training and testing data. However, in many real situations, such data is often unavailable [5, 21, 37, 48],

\*Corresponding authors: Li Shen and Chun Yuan

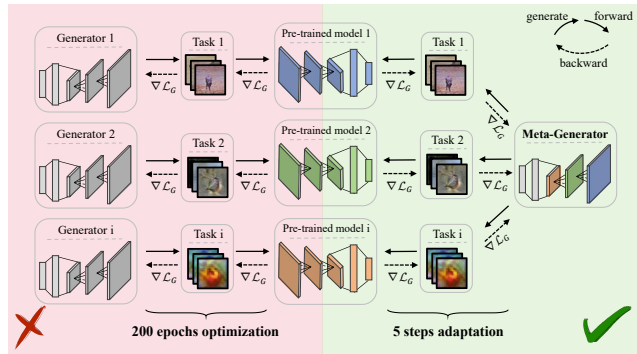


Figure 1. Faster Inversion via Meta-Generator significantly enhances the efficiency of task generation. Tasks recovered from pre-trained models are used for training in the data-free setting. For each task, prior works need to train a specific generator with hundreds of generate-forward-backward iterations, while we only need a 5-step adaptation using the single meta-generator.

primarily due to data privacy concerns, security risks, or usage rights. For example, numerous individuals and institutions release task-specific pre-trained models from diverse domains on platforms like GitHub or Hugging Face without training data released. Such real-world situations highlight the value of DFML: collect some pre-trained models with weaker generalization abilities, which likely originate from diverse domains online, and train a meta-learner with superior generalization ability for new tasks.

Existing DFML methods predominantly center on data recovery from pre-trained models. PURER [16] optimizes a learnable dataset through model inversion for each pre-trained model, subsequently sampling pseudo tasks for meta-learning. BiDf-MKD [17], on the other hand, trains generators using multiple black-box APIs, generating both the support and query sets distinctly for meta-learning.

In this work, we rethink DFML from two perspectives detailed in Sec. 3.2. Through this re-evaluation, we find two pivotal limitations in existing methods: (i) **time-consuming data recovery processes**, and (ii) **overlooking the heterogeneity among different pre-trained models**. Specifically,




Figure 2. Pre-trained models from different domains inherently exhibit distribution differences. Pre-trained models, even when trained on different classes of the same dataset, display variations in performance quality. As a result, their recovered tasks naturally present a gap in the distribution. Overlooking such model heterogeneity will cause the meta-learner to bias towards specific tasks, leading to local optima. Our proposed BELL optimizes the meta-learner by encouraging a positive inner product of gradients across tasks, thus enhancing its generalization ability.

prior research emphasizes generator learning *at the instance level*, where each pre-trained model is paired with a unique generator, designed to generate images consistent with specific training distribution. Training such a generator often necessitates hundreds of generate-forward-backward iterations, making the recovery process of training data considerably time-consuming. Moreover, due to the varying model architectures and qualities of different pre-trained models, some of which may have relatively low accuracy, and the possibility that they originate from distinct domains (see Fig. 2), their recovered tasks inherently present a distribution gap. We also present a *t*-SNE visualization to show the gap among them in Fig. 4. Applying simple Empirical Risk Minimization (ERM) might lead the meta-learner to be biased towards specific tasks, compromising its performance on others and limiting its generalization to unseen tasks.

Armed with these insights, we propose **Faster and Better Data-Free Meta-Learning (FREE)**, a unified framework that contains a meta-generator and a meta-learner. The meta-generator rapidly recovers specific tasks from pre-trained models for training the meta-learner, significantly accelerating the data recovery process. Concurrently, the meta-learner alleviates the gap among different tasks recovered from heterogeneous models, enhancing generalization to new unseen tasks. For **faster inversion via meta-generator (FIVE)**, we introduce the concept of a meta-generator learned *at the task level* (see Fig. 1). The meta-generator is trained across all pre-trained models for capturing shared representational knowledge. Drawing inspiration from meta-learning, we interpret *each pre-trained model as a distinct task*. The meta-generator is designed to yield minimal loss for each pre-trained model, adapting through only five steps using its self-generated data. The inner loop’s objective is to hone the rapid adaptability of the meta-generator, while the outer loop focuses on its generalization across multiple pre-trained models. For **better generalization via meta-learner (BELL)**,

we introduce an implicit gradient alignment algorithm across tasks from different pre-trained models, utilizing multi-task knowledge distillation to train the meta-learner. We treat the current task as the outer loop and its conflicting tasks as the inner loop, encouraging the optimization path to be suitable for all tasks. For instance, if the gradient directions of  $g_i$  and  $g_j$  are in alignment such that  $g_i \cdot g_j > 0$ , gradient descent updates will alleviate conflicts between different tasks [33]. By advocating a shared gradient direction, our meta-learner can extract common features across tasks/domains. This ensures that the meta-learner excels even on new unseen tasks, achieving superior generalization compared to ERM.

We conduct comprehensive experiments on three meta-learning benchmark datasets, *i.e.*, *mini-ImageNet*, *CIFAR-FS*, and *CUB*, for demonstrating our superiority over existing DFML methods. Compared to the state-of-the-art, our approach not only achieves a  $20\times$  speed-up but also shows an improvement (+ 1.42% ~ 4.78%). Furthermore, our approach effectively tackles the model heterogeneity in challenging multi-domain and multi-architecture scenarios.

In summary, our main contributions are three-fold:

- For the first time, we have a closer look at DFML, highlighting the pressing importance of addressing efficiency dilemma and model heterogeneity.
- To accelerate data recovery processes, we innovatively treat pre-trained models as tasks, focusing on a rapidly adaptive meta-generator. Recognizing the heterogeneity among different pre-trained models, we incorporate gradient alignment into the DFML framework. This alleviates inherent conflicts/gaps across tasks/domains, thereby improving the meta-learner’s generalization.
- Experiments on various benchmarks demonstrate the superiority of the proposed approach. Furthermore, we provide comprehensive discussions to elucidate the working mechanism of each component in our approach.

## 2. Related Work

**Meta-learning** [1, 3, 6, 18, 26, 28, 29, 35, 44], aka “learning-to-learn”, aims to acquire general prior knowledge from a large collection of tasks, enabling learning systems to rapidly adapt to novel categories using only a few examples. Further, Nichol et al. [24] introduce Reptile as an evolution of MAML [10]. This work underscores the link between the inner loop update and the maximization of the gradient inner product. Reptile aims to align gradients across minibatches from the same task, which enhances within-task generalization. In contrast, our approach formulates the inner loop using multiple tasks recovered from heterogeneous pre-trained models, which fosters across-task generalization.

For the heterogeneity in meta-learning, Yao et al. propose to cluster tasks into hierarchical structures enabling customized knowledge to heterogeneous tasks [42], or disentangle the meta-learner as a graph with different knowledge

blocks to learn from heterogeneous tasks [43].

**Data-free meta-learning (DFML)** [16, 17, 19, 40] is to ensure adaptation of acquired knowledge to new unseen tasks without training data. By leveraging multiple pre-trained models with weaker generalization abilities, which likely originate from diverse domains online, the intent is to learn a meta-learner with superior generalization ability. Regarding previous works, Wang et al. [40] suggest predicting the meta initialization using a neural network that operates in the parameter space. Recently, PURER [16] introduces a progressive method to synthesize a series of pseudo-tasks, inverting training data from each pre-trained model. BiDf-MKD [17] meta-learns the meta initialization by transferring meta knowledge from a collection of black-box APIs via zero-order gradient estimation. Fang et al. [9] propose to reuse common features across different images, rather than optimizing each image independently. In our study, we apply meta-learning principles at the pre-trained model level to accelerate the training process.

**Gradient alignment** [8, 20, 33] emphasizes the alignment or consistency of gradients during model training, and empirical evidence has shown its advantages in fields such as continual learning [39], federated learning [7], and multi-task learning [15]. Specifically, Ren et al. [31] focus on employing the gradient inner product to assign importance weights to training examples. PCGrad [47] introduces a gradient surgery technique that projects the gradient of one task onto the normal plane of the conflicting gradient of another task. Recently, Patel et al. [25] identify an implicit aligning factor between the knowledge-retention and knowledge-acquisition in data-free knowledge distillation. Nonetheless, we identify and address the heterogeneity that exists among different pre-trained models, introducing a more effective task-aligned meta-learning approach.

### 3. A Closer Look at DFML

In this section, we first review how DFML learns from a collection of pre-trained models. Then, we rethink DFML from two perspectives: efficiency and model heterogeneity.

#### 3.1. Preliminary

**DFML Setup.** We are given only a collection of pre-trained models  $\mathcal{M}_{pool} = \{M_i\}$ , each designed to solve a specific task without accessing their training data. DFML aims to learn a meta-learner  $F(\cdot; \theta)$  that can be rapidly transferred to new unseen tasks using pre-trained models  $\mathcal{M}_{pool}$ . During meta-testing, we sample “ $N$ -way  $K$ -shot” tasks. For such a task, there are  $N$  classes and each class only has  $K$  labeled samples  $(\mathcal{S}, \mathcal{Y}_{\mathcal{S}})$  named the support set, and  $U$  unlabelled samples per class  $Q$  called the query set, where  $\mathcal{Y}_Q$  is the corresponding ground truth. We use the support set  $(\mathcal{S}, \mathcal{Y}_{\mathcal{S}})$  to adapt the meta-learner  $F(\cdot; \theta)$  to each specific task. The query set  $Q$  is what we actually need to predict.

Table 1. GFLOPs of one image’s gradient back-propagation.

Resolution	Conv-4	ResNet-18	ResNet-50
$32 \times 32$	0.03	1.12	2.62
$84 \times 84$	0.23	7.87	18.36
$224 \times 224$	1.63	54.67	128.34

**Model inversion.** Model inversion [11, 22, 25, 46] aims to recover the training data  $\hat{X}$  from the pre-trained model  $M$  as an alternative to the inaccessible original data  $X$ . So a generator  $G(\cdot; \theta_G)$  can be introduced to estimate the distribution of  $X$  by designing an inversion loss with the pre-trained model  $M$ . Specifically, given a latent code  $z$  as a low-dimensional representation, the generator maps  $z$  to the intended data approximation  $\hat{x}$ . In this case, the generator is only responsible for a small part of the data distribution. Compared with the pixel updating strategy used in [16] that updates different pixels independently, the generator can provide stronger regularization on pixels because they are produced from the shared weights.

#### 3.2. Rethinking Existing DFML Methods

**Efficiency dilemma.** Model inversion generally requires a *generate-forward-backward* computation for each optimization step, with the gradient back-propagation process being the most time-intensive [14, 32]. As the structure of the pre-trained model expands, or the scale of the generated data increases, the time and computational costs rise substantially. We quantitatively measure the GFLOPs required for the gradient back-propagation of a single image through different model architectures, as shown in Tab. 1. The total computation for inversion can be expressed as  $N_{iter} \times N_{img} \times \text{GFLOPs}$ . Prior works optimize pixels individually or train a unique generator from scratch for each pre-trained model to reconstruct specific training distribution. Such pixel optimization or generator training often demands hundreds of generate-forward-backward computations. However, certain foundational features and textures are *shared* across different tasks, typically as the shallow layers of the generator [27]. If we train a meta-generator that retains common initialization parameters, rapid adaptation to specific tasks can be achieved in a few steps.

**Model heterogeneity.** Pre-trained models which contain random classes can be perceived as sub-domains with their diverse data distributions, due to statistical properties differences such as pixel intensity or texture variation [4, 36]. Diverse model architectures also contribute to model heterogeneity, influencing how data is represented. Moreover, our knowledge regarding the quality of pre-trained models is limited, as we have *no data* for evaluation. Given the chance that they stem from different domains, the tasks recovered from them inherently exhibit a distribution gap. We provide

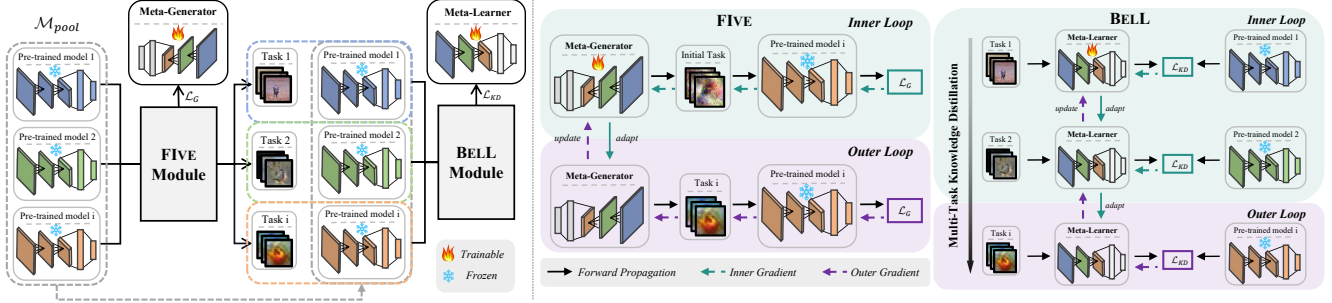


Figure 3. An illustration of the proposed DFML framework. The framework consists of multiple pre-trained models ( $\mathcal{M}_{pool}$ ), a meta-generator and a meta-learner. The model inversion loss ( $\mathcal{L}_G$ ) optimizes the meta-generator, while the knowledge distillation loss ( $\mathcal{L}_{KD}$ ) optimizes the meta-learner. After adapting to pre-trained models, the meta-generator recovers specific tasks. The meta-learner learns from recovered tasks and their respective pre-trained models by multi-task knowledge distillation.

a  $t$ -SNE visualization to illustrate the gap among them in Fig. 4. Since neural networks tend to favor shortcuts and frequently learn simpler biases [13], a straightforward ERM across different tasks will cause the meta-learner to develop an undue bias towards specific tasks, limiting its generalization capabilities for more challenging unseen tasks. As shown in Tab. 5, applying ERM leads to a 3.04% decrease in our approach’s performance on the unseen test set.

## 4. Methodology

Inspired by the above observation, we propose a unified framework FREE as illustrated in Fig. 3. In Sec. 4.1, we describe our proposed FIVE module for accelerating model inversion, followed by adaptive task recovery and meta-generator learning. In Sec. 4.2, we propose BELL module to align gradients of different tasks recovered from the heterogeneous model, training a meta-learner for new tasks.

### 4.1. Faster Inversion via Meta-Generator (FIVE)

Synthesizing training data from pre-trained models is a primary task in DFML. We propose viewing these pre-trained models as distinct *tasks* and employing the meta-learning inspired strategy to train a *meta-generator*. This meta-generator can rapidly adapt to a given pre-trained model within  $k$  steps, generating task-specific data for training the meta-learner, significantly reducing the time for model inversion. We initially discuss how to inverse data from pre-trained models using the meta-generator, followed by how to learn the meta-generator’s parameters.

**Adaptive Task Recovery.** The meta-generator  $G(\cdot; \theta_G)$  takes the standard Gaussian noise  $\mathbf{Z}$  as inputs and outputs the recovered data  $\hat{\mathbf{X}} = G(\mathbf{Z}; \theta_G)$ . For a pre-trained model  $M$ , we aim to recover a pseudo task  $\mathcal{T} = \{(\hat{\mathbf{X}}, \mathbf{Y})\}$  via the meta-generator. In order to obtain the training distribution of the pre-trained model, we need to perform  $k$  steps of rapid adaptation on the meta-generator. We utilize the DeepInversion loss [45] for updating  $\mathbf{Z}$  and  $\theta_G$ , which contains

a classification loss to condition the recovered data on pre-defined target label  $\mathbf{Y}$ , and an adversarial loss to encourage data diversity by maximizing the KL divergence between the pre-trained model  $M$  and the meta-learner  $F(\cdot; \theta)$ :

$$\mathcal{L}_G = CE(M(\hat{\mathbf{X}}), \mathbf{Y}) - \eta \cdot KL(M(\hat{\mathbf{X}}), F(\hat{\mathbf{X}}; \theta)), \quad (1)$$

where  $\eta = \mathbb{I}\{\arg \max M(\hat{\mathbf{X}}) = \arg \max F(\hat{\mathbf{X}}; \theta)\}$  and  $\hat{\mathbf{X}} = G(\mathbf{Z}; \theta_G)$ . The function  $\mathbb{I}(\cdot)$  is an indicator to enable  $\hat{\mathbf{X}}$  with the same prediction from the pre-trained model and the meta-learner ( $\eta = 1$ ), otherwise disable it ( $\eta = 0$ ).

For each pre-trained model  $M_i$ , we first clone the meta-generator  $G(\cdot; \theta_G)$  and randomly sample a latent code  $\mathbf{Z}$ . After a  $k$ -step adaptation as the inner loop, we obtain task-specific parameters  $(\mathbf{Z}^i, \theta_G^i)$ . Then, we can recover specific data  $\hat{\mathbf{X}}_i = G(\mathbf{Z}^i; \theta_G^i)$  from the pre-trained model  $M_i$ .

**Meta-Generator Learning.** In the inner loop, the adaptation process provides a sequence of losses  $\mathcal{L}_G^0, \mathcal{L}_G^1, \dots, \mathcal{L}_G^{k-1}$ , *i.e.*, losses on different data generated by the updating generator over  $k$  steps. To allow the meta-generator to build an internal representation suitable for a wide range of pre-trained models, the outer loop attempts to make the task-specific parameters reachable within the  $k$ -step adaptation:

$$\begin{aligned} \min_{\theta_G} \mathcal{L}_G^k(\hat{\mathbf{X}}_i) &= \mathcal{L}_G^k(G(\mathbf{Z}^i; \theta_G^i)), \\ \text{s.t. } (\mathbf{Z}^i, \theta_G^i) &= (\mathbf{Z}, \theta_G) - \nabla \mathcal{L}_G^0 - \nabla \mathcal{L}_G^1 - \dots - \nabla \mathcal{L}_G^{k-1}. \end{aligned} \quad (2)$$

When optimizing Eq. (2) via gradient descent, we also accelerate the meta-generator learning in back-propagation. We compute the gradient at  $\theta_G^i$  alternatively [10], and use the approximated gradient  $\nabla_{\theta_G^i} \mathcal{L}_G^k$  to update the meta-generator  $G(\cdot; \theta_G)$  as the outer loop (*cf.* Algorithm 1).

### 4.2. Better Generalization via Meta-Learner (BELL)

Pseudo tasks recovered from heterogeneous pre-trained models exhibit the distribution gap. To utilize them for jointly



---

**Algorithm 1: FIVE**

---

```
1 Input: Meta-generator  $G(\cdot; \theta_G)$ ; meta-learner  $F(\cdot; \theta)$ ;
   pre-trained models  $\{M_i\}_{i=1}^b$  in a batch.
2 Output: Optimized meta-generator  $G(\cdot; \theta_G)$ ;
   recovered tasks  $\{\mathcal{T}_i\}_{i=1}^b$  in a batch.
3 for index  $i \leftarrow 1$  to  $b$  do
4   Randomly sample  $\mathcal{Z}$  and clone  $G(\cdot; \theta_G)$ 
   // After a  $k$ -step adaptation
5   Obtain task-specific  $(\mathcal{Z}^i, \theta_G^i)$  from  $M_i$  w.r.t. Eq. (1)
6   Generate recovered data  $\hat{\mathcal{X}}_i = G(\mathcal{Z}^i; \theta_G^i)$ 
   // Meta-generator learning
7   Compute  $\mathbf{g}_G^i = \nabla_{\theta_G^i} \mathcal{L}_G^k(\hat{\mathcal{X}}_i)$  w.r.t. Eq. (2)
   // Outer loop for meta-generator
8 Update  $\theta_G \leftarrow \theta_G - \gamma \frac{1}{b} \sum_{i=1}^b \mathbf{g}_G^i$ 
```

---

training a meta-learner, we propose to align gradient directions of different tasks, facilitated by a multi-task knowledge distillation meta-learning algorithm.

**Multi-Task Knowledge Distillation.** Upon obtaining the recovered tasks from pre-trained models, we propose executing knowledge distillation to transfer the task-specific knowledge from the pre-trained model (acting as the teacher) to the meta-learner (acting as the student) using the recovered task. We optimize the meta-learner by minimizing the disagreement of predictions between them. This approach is favored because knowledge distillation offers richer supervision. It leverages the semantic class relationships present in the soft-label predictions from the pre-trained model  $M$ , as opposed to solely relying on the hard-label supervision from the generated data  $\hat{\mathcal{X}}$ .

For acquiring general knowledge from task-specific pre-trained models, we learn from multiple recovered tasks sampled from the pre-trained model pool  $\mathcal{M}_{pool}$  at the same time. A challenge arises when these tasks exhibit conflicting optimization directions (their gradients inner product  $< 0$ ). In such scenario, optimizing for task  $\mathcal{T}_i$  could inadvertently deteriorate the performance on task  $\mathcal{T}_j$  and vice versa. Unlike conventional meta-learning approaches which split a single task into a support set for the inner loop and a query set for the outer loop, we seek to align conflicting tasks by learning different tasks for the inner and outer loops. Specifically, we optimize the meta-learner  $F(\cdot; \theta)$  across recovered tasks in a sequence-by-sequence manner. In each sequence, the meta-learner is optimized for the task  $\mathcal{T}_i$ , while minimizing interference with other tasks:

$$\begin{aligned} \min_{\theta} \mathbb{E}_{\mathcal{T}_i \sim \mathcal{P}_M} \mathcal{L}_{KD}(\mathcal{T}_i; \tilde{\theta}) &\triangleq KL(M_i(\hat{\mathcal{X}}_i), F(\hat{\mathcal{X}}_i; \tilde{\theta})), \\ \text{s.t. } \tilde{\theta} &= \min_{\theta} \mathbb{E}_{\mathcal{T}_j \sim \mathcal{I}(\mathcal{T}_i)} \mathcal{L}_{KD}(\mathcal{T}_j; \theta), \end{aligned} \quad (3)$$

where  $\mathcal{P}_M$  represents the joint task distribution of pre-trained models  $\mathcal{M}_{pool}$ , and  $\mathcal{I}(\mathcal{T}_i)$  denotes the set of other tasks in the sequence that come before task  $\mathcal{T}_i$ .

---

**Algorithm 2: BELL**

---

```
1 Input: Meta-learner  $F(\cdot; \theta)$ ; recovered tasks  $\{\mathcal{T}_i\}_{i=1}^b$  and
   corresponding pre-trained models  $\{M_i\}_{i=1}^b$  in a batch.
2 Output: Optimized meta-learner  $F(\cdot; \theta)$ .
3 Clone meta-learner  $F(\cdot; \tilde{\theta}) \leftarrow F(\cdot; \theta)$ 
   // Multi-task knowledge distillation
4 for index  $i \leftarrow 1$  to  $b$  do
5   Compute  $\mathbf{g}_i = \nabla_{\tilde{\theta}} \mathcal{L}_{KD}(\mathcal{T}_i; \tilde{\theta})$  w.r.t. Eq. (3)
6   Update  $\tilde{\theta} \leftarrow \tilde{\theta} - \alpha \mathbf{g}_i$ 
   // Outer loop for meta-learner
7 Update  $\theta \leftarrow \theta + \epsilon(\tilde{\theta} - \theta)$ 
```

---

**Implicit Gradient Alignment.** To understand how it results in the desired alignment between different tasks, we can analyze the proposed objective in Eq. (3). We employ Taylor’s expansion to elucidate the connection between the gradient  $\nabla_{\tilde{\theta}} \mathcal{L}_{KD}(\mathcal{T}_i; \tilde{\theta})$  and gradient at initial point  $\nabla_{\theta} \mathcal{L}_{KD}(\mathcal{T}_i; \theta)$ , as described in Lemma 1.

**Lemma 1.** *If  $\mathcal{L}_{KD}$  has Lipschitz Hessian, then:*

$$\begin{aligned} \nabla_{\tilde{\theta}} \mathcal{L}_{KD}(\mathcal{T}_i; \tilde{\theta}) &= \nabla_{\theta} \mathcal{L}_{KD}(\mathcal{T}_i; \theta) + O(\alpha^2) \\ &\quad - \alpha \nabla_{\tilde{\theta}}^2 \mathcal{L}_{KD}(\mathcal{T}_i; \theta) \nabla_{\theta} \mathcal{L}_{KD}(\mathcal{T}_j; \theta), \end{aligned}$$

where  $\alpha$  is the step size of the inner loop.

**Theorem 1.** *If  $\mathcal{T}_i$  can be regarded as independent identically distributed samples from the distribution  $\mathcal{P}_M$ , then:*

$$\begin{aligned} \nabla_{\theta} \mathcal{L}_{KD}(\mathcal{T}_i; \tilde{\theta}) &= \nabla_{\theta} \mathcal{L}_{KD}(\mathcal{T}_i; \theta) + O(\alpha^2) \\ &\quad - \alpha \nabla \underbrace{(\nabla_{\theta} \mathcal{L}_{KD}(\mathcal{T}_i; \theta) \nabla_{\theta} \mathcal{L}_{KD}(\mathcal{T}_j; \theta))}_{\text{Gradient Alignment}}, \end{aligned}$$

i.e., the inner product between gradients of different tasks.

From the analysis above, we observe that the gradient of  $\mathcal{T}_i$  produces the inner product with other tasks that might pose conflicts. This indicates, optimizing the meta-learner minimizes the expected loss over tasks  $\mathcal{T}_i \sim \mathcal{P}_M$  (effectively the ERM), and maximizes the inner product between gradients of different tasks  $\nabla_{\theta} \mathcal{L}_{KD}(\mathcal{T}_i; \theta) \nabla_{\theta} \mathcal{L}_{KD}(\mathcal{T}_j; \theta)$ . Hence, it enforces the meta-learner to seek a common direction, encouraging across-task generalization.

**Optimization.** However, optimizing the meta-learner based on Eq. (3) requires high-order gradients when computing  $\nabla_{\theta} \mathcal{L}_{KD}(\mathcal{T}_i; \tilde{\theta})$ , making the back-propagation highly inefficient. We could apply a first-order approximation [24] to further accelerate the gradient computing. We continue to update  $\tilde{\theta}$  using the gradient  $\nabla_{\tilde{\theta}} \mathcal{L}_{KD}(\mathcal{T}_i; \tilde{\theta})$ . Then, the meta-learner  $F(\cdot; \theta)$  is updated through a weighted difference between the resulted parameters  $\tilde{\theta}$  and the original parameters  $\theta$  (cf. Algorithm 2).

---

**Algorithm 3:** FREE

---

1 **Input:** A collection of pre-trained models  $\mathcal{M}_{pool}$ ;  
a meta-generator  $G(\cdot; \theta_G)$ ; the memory bank  $\mathcal{B}$ .  
2 **Output:** A meta-learner  $F(\cdot; \theta)$  for new unseen tasks.  
3 **for** epoch  $i \leftarrow 1$  **to**  $N$  **do**  
4     Sample pre-trained models  $\{M_i\}_{i=1}^b$  from  $\mathcal{M}_{pool}$   
5     Recover specific tasks  $\{\mathcal{T}_i\}_{i=1}^b$  with  $G(\cdot; \theta_G)$   
6     Put into memory bank  $\mathcal{B} \leftarrow \mathcal{B} + \{\mathcal{T}_i\}_{i=1}^b$   
7     Update meta-generator  $\theta_G$  w.r.t. Eq. (2)  
8     Update meta-learner  $\theta$  with  $\{M_i\}_{i=1}^b$  and  $\{\mathcal{T}_i\}_{i=1}^b$   
9     Construct interpolated tasks  $\hat{\mathcal{T}}$  from  $\mathcal{B}$   
10    Update meta-learner  $\theta$  w.r.t. Eq. (4)

---

**Cross Task Replay.** A small number of pre-trained models (e.g., 100) are insufficient to represent the actual underlying task distribution, leading to an over-reliance on specific tasks. Hence, we employ a memory bank  $\mathcal{B}$  with a first-in-first-out structure to store the previously recovered tasks, and replay new tasks interpolated across different tasks, i.e., a random combination of classes. This further alleviates potential conflicts between tasks, assisting the meta-learner in generalizing across tasks. For these replayed tasks  $\hat{\mathcal{T}}$ , which include the support set  $\mathcal{S}$ , the query set  $\mathcal{Q}$  and class labels  $\mathcal{Y}$ , we update the meta-learner as follows:

$$\begin{aligned} \min_{\theta} \mathbb{E}_{\hat{\mathcal{T}} \sim \mathcal{P}_{\mathcal{B}}} \mathcal{L}_{outer} &\triangleq \mathcal{L}_{CE}(F(\mathcal{Q}; \theta_c), \mathcal{Y}_{\mathcal{Q}}), & (4) \\ \text{s.t. } \theta_c = \min_{\theta} \mathcal{L}_{inner} &\triangleq \mathcal{L}_{CE}(F(\mathcal{S}; \theta), \mathcal{Y}_{\mathcal{S}}). \end{aligned}$$

In the end, we summarize the main pipeline of the proposed framework in Algorithm 3.

## 5. Experiments

In this section, we begin by detailing our experimental setup, including an overview of the compared baselines. Following that, we present the main results of our approach, evaluating performance based on meta-testing accuracy and meta-training speed. Additionally, we offer ablation studies and discussions to facilitate comprehensive analyses.

### 5.1. Experimental Setup

**Datasets and pre-trained models.** We conduct experiments on two widely-used DFML benchmark datasets, and one fine-grained dataset, including *miniImageNet* [30], *CIFAR-FS* [2] and *CUB* [38]. Following standard splits [41], we split each dataset into the meta-training, meta-validating and meta-testing subsets with disjoint label spaces. In the DFML setting, we have no access to the meta-training data. Following [17, 40], we collect 100 models pre-trained on 100  $N$ -way tasks sampled from the meta-training subset, and those models are used as the meta-training resources.

**Implementation details.** For the model architecture, we adopt Conv4 as the architecture of the meta-learner and the pre-trained models for a fair comparison with existing works. We provide the detailed structure for the meta-generator in Appendix B. For hyperparameters, the batch size  $b$  is set to 4, and the learning rate  $\gamma$ ,  $\alpha$ , and  $\epsilon$  are all set to 0.001. We report the average accuracy over 600 meta-testing tasks. We leave the other setup in Appendix A.

**Baselines.** (i) **Random.** Learn a classifier using the support set from scratch for each meta-testing task. (ii) **Best-Model.** We select the pre-trained model with the highest reported accuracy to directly predict the query set during meta-testing. (iii) **Average.** Average all pre-trained models and then fine-tune it using the support set. (iv) **OTA** [34]. Calculate the weighted average of all pre-trained models and then fine-tune it using the support set. (v) **DRO** [40]. Meta-learn a hyper-network to fuse all pre-trained models into one single model, which serves as the meta-initialization and can be adapted to each meta-testing task using the support set. (vi) **PURER** [16]. Adversarially train the meta-learner with a learnable dataset, where a batch of pseudo tasks is sampled for meta-training at each iteration. (vii) **BiDf-MKD** [17]. A bi-level data-free meta knowledge distillation framework to transfer general knowledge in the white-box setting.

**Metrics.** In addition to the standard few-shot classification accuracy, we also focus on the speed of model inversion in DFML, quantifying this through the metric of GPU hours. Note that the time cost of meta-learner training is omitted since we only focus on the data recovery process. For consistency in comparison, all GPU hours are obtained using a single GeForce RTX 3090 GPU.

### 5.2. Main Results

**Comparisons with baselines.** Tab. 2 shows the results for 5-way classification compared with current baselines. To ascertain the efficacy of FREE, we compare it against two DFML algorithm categories: the non-inversion methods, which facilitate model fusion in the parameter space, and the inversion-based methods that generate pseudo data for training. As shown in Tab. 2, our approach considerably outperforms all non-inversion methods, and is significantly faster than other inversion-based methods. For instance, PURER generates data by constructing a smaller pseudo dataset, with each class comprising only 20 images. Nevertheless, PURER demands 15000 iterations to converge in optimizing the dataset. In contrast, we only require 2/5 steps for different tasks, owing to our adaptive meta-generator.

Beyond speed enhancements, FREE also achieves better accuracy across all three datasets. Compared with the foremost baseline BiDf-MKD, FREE displays performance enhancements ranging from 0.32%~2.37% for 1-shot learning and from 1.42%~4.78% for 5-shot learning. These improvements can be attributed to two key factors. Firstly,

Table 2. Compare to existing baselines in DFML. Time: GPU hours taken by the data recovery process, so non-inversion methods only report accuracy.  $\text{FREE}_{2/5}$  denotes the 2-step and 5-step adaptation of the meta-generator, respectively.

Method	CIFAR-FS			miniImageNet			CUB		
	5-way 1-shot	5-way 5-shot	Time	5-way 1-shot	5-way 5-shot	Time	5-way 1-shot	5-way 5-shot	Time
Random	28.59 ± 0.56	34.77 ± 0.62	-	25.06 ± 0.50	28.10 ± 0.52	-	26.26 ± 0.48	29.89 ± 0.55	-
Best-Model	21.68 ± 0.66	25.05 ± 0.67	-	22.86 ± 0.61	26.26 ± 0.63	-	24.16 ± 0.73	29.16 ± 0.73	-
Average	23.96 ± 0.53	27.04 ± 0.51	-	23.79 ± 0.48	27.49 ± 0.50	-	24.53 ± 0.46	28.00 ± 0.47	-
OTA [34]	29.10 ± 0.65	34.33 ± 0.67	-	24.22 ± 0.53	27.22 ± 0.59	-	24.23 ± 0.60	25.42 ± 0.63	-
DRO [40]	30.43 ± 0.43	36.21 ± 0.51	-	27.56 ± 0.48	30.19 ± 0.43	-	28.33 ± 0.69	31.24 ± 0.76	-
PURER [16]	38.66 ± 0.78	51.95 ± 0.79	1.21h	31.14 ± 0.63	40.86 ± 0.64	1.31h	30.08 ± 0.59	40.93 ± 0.66	1.97h
BiDf-MKD [17]	37.66 ± 0.75	51.16 ± 0.79	2.47h	30.66 ± 0.59	42.30 ± 0.64	8.87h	31.62 ± 0.60	44.32 ± 0.69	7.16h
<b>FREE<sub>2</sub></b>	36.56 ± 0.73	47.31 ± 0.76	0.05h	30.06 ± 0.61	41.60 ± 0.68	0.25h	30.64 ± 0.58	39.43 ± 0.61	0.20h
<b>FREE<sub>5</sub></b>	<b>39.13 ± 0.85</b>	<b>52.58 ± 0.77</b>	0.11h	<b>33.03 ± 0.69</b>	<b>45.45 ± 0.69</b>	0.41h	<b>31.94 ± 0.61</b>	<b>49.10 ± 0.68</b>	0.32h

Table 3. Compare to baselines in a multi-domain scenario.

Method	CIFAR-FS + miniImageNet + CUB		
	5-way 1-shot	5-way 5-shot	Time
Random	24.85 ± 0.54	28.35 ± 0.61	-
PURER [16]	29.67 ± 0.59	37.96 ± 0.60	3.43h
BiDf-MKD [17]	31.44 ± 0.64	40.96 ± 0.59	7.78h
<b>FREE<sub>2</sub></b>	30.39 ± 0.56	40.69 ± 0.68	0.20h
<b>FREE<sub>5</sub></b>	<b>31.51 ± 0.63</b>	<b>44.04 ± 0.66</b>	0.34h

our proposed BELL effectively aligns conflicting gradients across different tasks, directly improving the generalization capabilities. Secondly, the efficacy of our meta-generator is evident as it rapidly absorbs knowledge from pre-trained models and provides valuable training data.

**Multi-domain scenario.** We conduct experiments in a challenging multi-domain scenario where all pre-trained models are tailored to address distinct tasks from multiple meta-training subsets (*CIFAR-FS*, *miniImageNet*, and *CUB*). For meta-testing, we assess the meta-learner on unseen tasks spanning *CIFAR-FS*, *miniImageNet*, and *CUB* concurrently, demanding the meta-learner to possess generalization capabilities across multiple domains. Tab. 3 showcases our results. Our approach achieves commendable results, outperforming the baseline by 6.66% and 15.69% in 1-shot and 5-shot learning, respectively. This solidly demonstrates that BELL effectively aligns tasks from different domains, optimizing them towards a unified direction, resulting in robust generalization across various domains.

**Multi-architecture scenario.** We also conduct experiments in a multi-architecture scenario where each pre-trained model differs in architecture. For each pre-trained model, the architecture is randomly chosen from Conv4, ResNet-10 and ResNet-18. The results, presented in Tab. 4, demonstrate the efficacy of our approach. It outperforms all baselines and can apply to multi-architecture scenario without any

Table 4. Compare to baselines in a multi-architecture scenario.

Method	CIFAR-FS		
	5-way 1-shot	5-way 5-shot	Time
PURER [16]	39.15 ± 0.70	49.08 ± 0.74	1.75h
BiDf-MKD [17]	38.08 ± 0.80	50.58 ± 0.81	4.33h
<b>FREE<sub>5</sub></b>	<b>39.63 ± 0.82</b>	<b>52.12 ± 0.79</b>	0.13h

Table 5. Effects of the proposed modules.

Setting	Module		CIFAR-FS	
	FIVE	BELL	5-way 1-shot	5-way 5-shot
Baseline			36.33 ± 0.81	47.67 ± 0.75
ERM	✓		37.43 ± 0.77	49.54 ± 0.73
Sequence		✓	38.56 ± 0.78	50.78 ± 0.79
<b>FREE<sub>5</sub></b>	✓	✓	<b>39.13 ± 0.85</b>	<b>52.58 ± 0.77</b>

change. This flexibility is attributable to our method’s lack of constraints on the underlying architecture or scale of the pre-trained models, which enables effective task alignment across heterogeneous models.

### 5.3. Discussions

**Ablation studies.** Tab. 5 evaluates the impact of each module on the *CIFAR-FS*. Initially, we set a baseline only performing meta-learning via cross task replay, using a sequential generator. The generator trained on a previous pre-trained model is directly used as an initialization for subsequent tasks. However, updating the generator with only 5 steps for each pre-trained model proves to be insufficient. This is attributed to the non-reusability of the sequential generator across diverse pre-trained models. Incorporating FIVE module offers notable enhancements, suggesting that the meta-generator can rapidly adapt and inverse task-specific training data. Nonetheless, the current method (*i.e.*, ERM) still merely aggregates knowledge distillation losses from multiple tasks to update the meta-learner, overlooking the

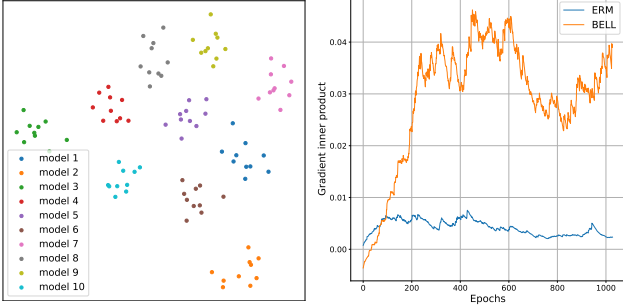


Figure 4. The  $t$ -SNE visualization of same samples input into different pre-trained models.

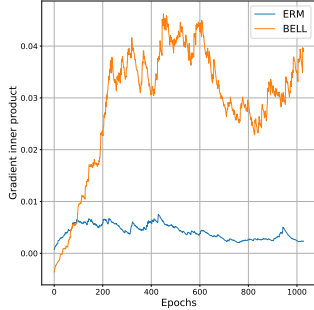


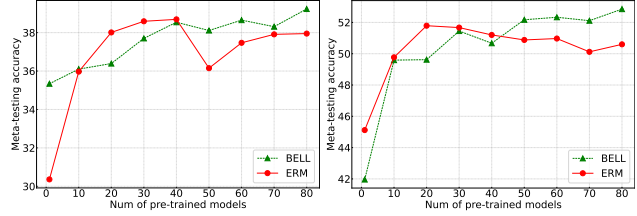
Figure 5. The gradient inner product across tasks from different pre-trained models.

alignment of optimization directions across tasks. By adding BELL module, we observe a significant improvement (*i.e.*, Sequence), demonstrating the effectiveness of the conflicting gradient alignment. With all modules, we achieve the best performance with a boosting improvement of 2.8% and 4.91%, showing their complementarity.

**Heterogeneity among pre-trained models.** We present a  $t$ -SNE visualization of the features extracted from different pre-trained models in Fig. 4. We randomly sample 10 images of the same class and use 10 pre-trained models from *CIFAR-FS* to extract their features. The visualization clearly shows that the features of the same image, when extracted by different pre-trained models, are scattered in the feature space. This implies that the inversed tasks from distinct pre-trained models naturally inherit varying feature distributions. This observation underscores the importance of addressing the heterogeneity among pre-trained models.

**Gradient inner product maximization.** In our approach, we incorporate BELL to align gradients across different tasks, guiding the meta-learner towards a unified direction. To further validate whether BELL implicitly maximizes the gradient inner product, we chart its evolution during meta-learner training in Fig. 5. We contrast BELL with the ERM method, monitoring the mean gradient inner product between tasks derived from varied pre-trained models in each inner loop. As evident from Fig. 5, the gradient inner product for BELL gradually rises with training progression. Conversely, ERM maintains lower values throughout. This confirms BELL’s effectiveness in implicitly aligning the gradient direction by enhancing the gradient inner product.

**Effects of the number of pre-trained models.** We conduct experiments on DFML with varying numbers of pre-trained models. Each model is pre-trained on a 5-way subset of *CIFAR-FS*. As shown in Fig. 6, the performance improves with an increasing number of models. This can be attributed to the fact that a greater quantity of pre-trained models offers a more comprehensive understanding of various classes, thereby enhancing generalization capabilities for unseen classes. Interestingly, we observe that ERM delivers



(a) 5-way 1-shot (b) 5-way 5-shot  
Figure 6. Effects of the number of pre-trained models.



Figure 7. Visualizations of the 5-way recovered task via our meta-generator. Each group represents one class.

promising results with a smaller set of pre-trained models, at times even outperforming our approach. However, as the number of pre-trained models continues to climb, ERM’s performance starts to drop and is eventually surpassed by our approach. This decline can be attributed to the growing issues of model heterogeneity and task conflicts with an increase in the number of pre-trained models. The poor performance of ERM highlights the challenge of utilizing numerous available pre-trained models from sources like GitHub to enable efficient learning on unseen downstream tasks. Contrarily, we effectively navigate this challenge, consistently delivering performance enhancements.

**Visualizations.** Fig. 7 presents a 5-way task recovered from the Conv4. A significant advantage of DFML lies in its ability to learn from weaker pre-trained models without data privacy leakage. Our approach generates images that are visually distinct from the originals, thereby addressing privacy concerns. The images we generate display high-frequency patterns and textures, which are recognized for their particular value in training neural networks [12, 23].

## 6. Conclusion

In this work, we highlight the necessity of addressing the efficiency dilemma and the heterogeneity among pre-trained models in DFML. To address the slow recovery speed, we train a meta-generator capable of rapidly adapting to specific tasks. To further address the potential conflicts in recovered tasks, we train a meta-learner to align the gradients of different tasks. This meta-learner captures task-invariant features, thus enabling generalization to new unseen tasks. Extensive experiments on multiple benchmarks show significant speed and performance gains in our approach.

**Acknowledgement.** This work is supported by the National Key R&D Program of China (2022YFB4701400/4701402), SSTIC Grant(KJZD20230923115106012), Shenzhen Key Laboratory (ZDSYS20210623092001004), and Beijing Key Lab of Networked Multimedia.



## References

- [1] Sungyong Baik, Janghoon Choi, Heewon Kim, Dohee Cho, Jaesik Min, and Kyoung Mu Lee. Meta-learning with task-adaptive loss function for few-shot learning. In *ICCV*, pages 9465–9474, 2021. 2
- [2] Luca Bertinetto, Joao F. Henriques, Philip Torr, and Andrea Vedaldi. Meta-learning with differentiable closed-form solvers. In *ICLR*, 2019. 6
- [3] Arnav Chavan, Rishabh Tiwari, Udbhav Bamba, and Deepak K Gupta. Dynamic kernel selection for improved generalization and memory efficiency in meta-learning. In *CVPR*, pages 9851–9860, 2022. 2
- [4] Binghui Chen, Zhaoyi Yan, Ke Li, Pengyu Li, Biao Wang, Wangmeng Zuo, and Lei Zhang. Variational attention: Propagating domain-specific knowledge for multi-domain learning in crowd counting. In *ICCV*, pages 16065–16075, 2021. 3
- [5] Hanting Chen, Yunhe Wang, Chang Xu, Zhaohui Yang, Chuanjian Liu, Boxin Shi, Chunjing Xu, Chao Xu, and Qi Tian. Data-free learning of student networks. In *ICCV*, pages 3514–3522, 2019. 1
- [6] Yinbo Chen, Zhuang Liu, Huijuan Xu, Trevor Darrell, and Xiaolong Wang. Meta-baseline: Exploring simple meta-learning for few-shot learning. In *ICCV*, pages 9062–9071, 2021. 2
- [7] Yatin Dandi, Luis Barba, and Martin Jaggi. Implicit gradient alignment in distributed and federated learning. In *AAAI*, pages 6454–6462, 2022. 3
- [8] Amir Erfan Eshratifar, David Eigen, and Massoud Pedram. Gradient agreement as an optimization objective for meta-learning. *arXiv preprint arXiv:1810.08178*, 2018. 3
- [9] Gongfan Fang, Kanya Mo, Xinchao Wang, Jie Song, Shitao Bei, Haofei Zhang, and Mingli Song. Up to 100x faster data-free knowledge distillation. In *AAAI*, pages 6597–6604, 2022. 3
- [10] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, pages 1126–1135, 2017. 2, 4
- [11] Ahmed Frikha, Haokun Chen, Denis Kropf, Thomas Runzler, and Volker Tresp. Towards data-free domain generalization. In *ACML*, pages 327–342, 2023. 3
- [12] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A Wichmann, and Wieland Brendel. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. In *ICLR*, 2018. 8
- [13] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2:665–673, 2020. 4
- [14] Audrunas Gruslys, Rémi Munos, Ivo Danihelka, Marc Lanctot, and Alex Graves. Memory-efficient backpropagation through time. In *NeurIPS*, 2016. 3
- [15] Shi Guanguyan, Qimai Li, Wenlong Zhang, Jiabin Chen, and Xiao-Ming Wu. Recon: Reducing conflicting gradients from the root for multi-task learning. In *ICLR*, 2022. 3
- [16] Zixuan Hu, Li Shen, Zhenyi Wang, Tongliang Liu, Chun Yuan, and Dacheng Tao. Architecture, dataset and model-scale agnostic data-free meta-learning. In *CVPR*, pages 7736–7745, 2023. 1, 3, 6, 7
- [17] Zixuan Hu et al. Learning to learn from apis:black-box data-free meta-learning. In *ICML*, 2023. 1, 3, 6, 7
- [18] Muhammad Abdullah Jamal and Guo-Jun Qi. Task agnostic meta-learning for few-shot learning. In *CVPR*, pages 11719–11727, 2019. 2
- [19] Nameyong Kwon, Hwidong Na, Gabriel Huang, and Simon Lacoste-Julien. Repurposing pretrained models for robust out-of-domain few-shot learning. In *ICLR*, 2020. 1, 3
- [20] Sanghyuk Lee, Seunghyun Lee, and Byung Cheol Song. Contextual gradient scaling for few-shot learning. In *WACV*, pages 834–843, 2022. 3
- [21] Weishi Li, Yong Peng, Miao Zhang, Liang Ding, Han Hu, and Li Shen. Deep model fusion: A survey. *arXiv preprint arXiv:2309.15698*, 2023. 1
- [22] Yuang Liu, Wei Zhang, Jun Wang, and Jianyong Wang. Data-free knowledge transfer: A survey. *arXiv preprint arXiv:2112.15278*, 2021. 3
- [23] Paul Micaelli and Amos J Storkey. Zero-shot knowledge transfer via adversarial belief matching. In *NeurIPS*, 2019. 8
- [24] Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018. 2, 5
- [25] Gaurav Patel, Konda Reddy Mopuri, and Qiang Qiu. Learning to retain while acquiring: Combating distribution-shift in adversarial data-free knowledge distillation. In *CVPR*, pages 7786–7794, 2023. 3
- [26] Xiaorong Qin, Xinhang Song, and Shuqiang Jiang. Bi-level meta-learning for few-shot domain generalization. In *CVPR*, pages 15900–15910, 2023. 2
- [27] Aniruddh Raghu, Maithra Raghu, Samy Bengio, and Oriol Vinyals. Rapid learning or feature reuse? towards understanding the effectiveness of maml. In *ICLR*, 2019. 3
- [28] Jathushan Rajasegaran, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Mubarak Shah. itaml: An incremental task-agnostic meta-learning approach. In *CVPR*, pages 13588–13597, 2020. 2
- [29] Aravind Rajeswaran, Chelsea Finn, Sham M Kakade, and Sergey Levine. Meta-learning with implicit gradients. In *NeurIPS*, 2019. 2
- [30] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *ICLR*, 2017. 6
- [31] Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. In *ICML*, pages 4334–4343, 2018. 3
- [32] Sunandini Sanyal, Sravanti Addepalli, and R Venkatesh Babu. Towards data-free model stealing in a hard label setting. In *CVPR*, pages 15284–15293, 2022. 3
- [33] Yuge Shi, Jeffrey Seely, Philip Torr, N Siddharth, Awni Hanun, Nicolas Usunier, and Gabriel Synnaeve. Gradient matching for domain generalization. In *ICLR*, 2021. 2, 3
- [34] Sidak Pal Singh and Martin Jaggi. Model fusion via optimal transport. In *NeurIPS*, pages 22045–22055, 2020. 6, 7
- [35] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *NeurIPS*, pages 4077–4087, 2017. 2

- [36] Tiecheng Song, Jie Feng, Lin Luo, Chenqiang Gao, and Hongliang Li. Robust texture description using local grouped order pattern and non-local binary pattern. *IEEE TCSVT*, 31: 189–202, 2020. 3
- [37] Jean-Baptiste Truong, Pratyush Maini, Robert J Walls, and Nicolas Papernot. Data-free model extraction. In *CVPR*, pages 4771–4780, 2021. 1
- [38] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The Caltech-UCSD birds-200-2011 dataset. Technical report, California Institute of Technology, 2011. 6
- [39] Zhenyi Wang, Li Shen, Tiehang Duan, Donglin Zhan, Le Fang, and Mingchen Gao. Learning to learn and remember super long multi-domain task sequence. In *CVPR*, pages 7982–7992, 2022. 3
- [40] Zhenyi Wang, Xiaoyang Wang, Li Shen, Qiuling Suo, Kaiqiang Song, Dong Yu, Yan Shen, and Mingchen Gao. Meta-learning without data via wasserstein distributionally-robust model fusion. In *UAI*, pages 2045–2055, 2022. 1, 3, 6, 7
- [41] Davis Wertheimer, Luming Tang, and Bharath Hariharan. Few-shot classification with feature map reconstruction networks. In *CVPR*, pages 8012–8021, 2021. 6
- [42] Huaxiu Yao, Ying Wei, Junzhou Huang, and Zhenhui Li. Hierarchically structured meta-learning. In *ICML*, pages 7045–7054, 2019. 2
- [43] Huaxiu Yao, Yingbo Zhou, Mehrdad Mahdavi, Zhenhui Li, Richard Socher, and Caiming Xiong. Online structured meta-learning. In *NeurIPS*, 2020. 3
- [44] Huaxiu Yao, Yu Wang, Ying Wei, Peilin Zhao, Mehrdad Mahdavi, Defu Lian, and Chelsea Finn. Meta-learning with an adaptive task scheduler. In *NeurIPS*, pages 7497–7509, 2021. 2
- [45] Hongxu Yin, Pavlo Molchanov, Jose M Alvarez, Zhizhong Li, Arun Mallya, Derek Hoiem, Niraj K Jha, and Jan Kautz. Dreaming to distill: Data-free knowledge transfer via deepinversion. In *CVPR*, pages 8715–8724, 2020. 4
- [46] Shikang Yu, Jiachen Chen, Hu Han, and Shuqiang Jiang. Data-free knowledge distillation via feature exchange and activation region constraint. In *CVPR*, pages 24266–24275, 2023. 3
- [47] Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. Gradient surgery for multi-task learning. In *NeurIPS*, pages 5824–5836, 2020. 3
- [48] Hongling Zheng, Li Shen, Anke Tang, Yong Luo, Han Hu, Bo Du, and Dacheng Tao. Learn from model beyond fine-tuning: A survey. *arXiv preprint arXiv:2310.08184*, 2023. 1